

μ -tempered metadynamics: Artifact independent convergence times for wide hills

Bradley M. Dickson^{1,*}

¹ Center for Epigenetics, Van Andel Research Institute, Grand Rapids Michigan 49503, USA.

(Dated: November 3, 2015)

Recent analysis of well-tempered metadynamics (WTmetaD) showed that it converges without mollification artifacts in the bias potential. Here we explore how metadynamics heals mollification artifacts, how healing impacts convergence time, and whether alternative temperings may be used to improve efficiency. We introduce “ μ -tempered” metadynamics as a simple tempering scheme, inspired by a related mollified adaptive biasing potential (mABP), that results in artifact independent convergence of the free energy estimate. We use a toy model to examine the role of artifacts in WTmetaD and solvated alanine dipeptide to compare the well-tempered and μ -tempered frameworks demonstrating fast convergence for hill widths as large as 60° for μ TmetaD.

I. INTRODUCTION

Metadynamics was proposed in 2002[3] as a method for escaping metastabilities, and recently a proof of convergence has emerged[4] for tempered versions, which speaks directly to the form metadynamics takes when it is implemented. A critical feature of that proof, which has driven this report, was that WTmetaD converges without mollification artifacts in the bias potential. The metadynamics literature has not been clear on this point, with discussion of “blurring”[5] and even demonstrations of failed convergence on practical timescales where “blurring” is observed[6]. It is an intriguing insight to find that WTmetaD, and the metadynamics family, converges without mollification artifact as it means the hill width does not impact the accuracy of free energy estimation.

This insight was not predicted in Ref. 1, where we derived a novel bias identical to the motivating expression for well-tempered metadynamics[2]. In Ref. 1 we showed that this motivating expression defines a bias potential which converges to an end state that includes mollification artifacts. We also showed how to estimate the exact free energy from the bias potential via deconvolution.

Here, we start out by looking at the disconnect between mABP and WTmetaD. The two methods have identical motivating expressions for the bias potential, and agree when the hill width is small, yet diverge from one another when the hill width is large: WTmetaD must heal artifacts in the bias to converge while the mABP bias converges independently of artifacts. Below it is estimated that the timescale for healing mollification artifacts scales exponentially with the hill width for WTmetaD. Motivated to find a tempering schedule that improves upon the expensive healing process, we generalize the tempering of metadynamics into a common framework and propose an alternative schedule. The new scheme, called μ -tempering, produces a free energy estimate that is independent of the bias potential once the hills are small enough to allow the dynamics to be near equilibrium.

In other words, artifacts do not need to be healed. We demonstrate that this tempering overcomes the slow convergence of WTmetaD when α is large, allowing use of aggressive filling rates.

II. RESULTS

A. Metadynamics from mABP

The point of this section is to point out the disconnect between mABP and WTmetaD, and to restate a main result of mABP, which was that accurate free energy estimates could be made while employing arbitrary hill widths and one need not spend any time healing artifacts in the bias potential.

We first define some notation and briefly revisit prior results. We write \mathbf{x} for a single configuration in the n -dimensional configuration space \mathcal{X} of a dynamical system. Assume N collective variables (CVs) are specified and, of course, let these CVs be a good representation of \mathcal{X} . Ω is the space of CVs and $\boldsymbol{\xi}$ is a point in Ω . We define the free energy estimate (up to a constant ζ)

$$\zeta e^{-\beta A_\alpha(\boldsymbol{\xi})} = Z^{-1} \int_{\mathcal{X}} \delta_\alpha(\boldsymbol{\xi}(\mathbf{x}) - \boldsymbol{\xi}) e^{-\beta V(\mathbf{x})} d\mathbf{x}, \quad (1)$$

where

$$\delta_\alpha(\boldsymbol{\xi}) = \exp\left(-\frac{|\boldsymbol{\xi}|^2}{\alpha^2}\right) \quad (2)$$

and $\beta = k_B T$. We let Z take the normalization of δ_α .

In reference 1 we derived the following biasing potential from an intuitive on-the-fly reweighting scheme

$$V_b(\boldsymbol{\xi}, t) = \beta^{-1} \frac{b}{1-b} \ln[c(1-b) h_\alpha(\boldsymbol{\xi}, t) + 1] \quad (3)$$

where

$$h_\alpha(\boldsymbol{\xi}, t) = \int_0^t \delta_\alpha(\boldsymbol{\xi}(\mathbf{x}_s) - \boldsymbol{\xi}) ds \quad (4)$$

*Electronic address: bradley.dickson@vai.org

This biasing potential was designed to facilitate estimation of A_α by enhancing exploration of configuration space. We elaborated that the potential of mean force (PMF) in the biased ensemble is related to the bias potential as $V_b = -b\mu_\alpha/(1-b)$ where at long times $\mu_\alpha = -\beta^{-1}\ln[h_\alpha] = -\beta^{-1}\ln[\delta_\alpha * h]$. h is a histogram with vanishingly small α , and $*$ indicates a convolution. Thus the exact PMF is $\mu = -\beta^{-1}\ln[h] = A + V_b$. We showed that the free energy could be estimated from these dynamics as

$$A(\xi) = \mu(\xi) + \frac{b}{1-b}\mu_\alpha(\xi) \quad (5)$$

When one can take $\mu_\alpha \approx \mu$ the free energy estimate becomes

$$A(\xi, t) = -\frac{1}{b}V_b(\xi, t) \quad (6)$$

suggesting that $100 \times b$ percent of the free energy is cancelled by the bias potential. When α is so large that μ_α cannot replace μ ,

$$A(\xi, t) = -\beta^{-1}\ln[\delta_\alpha^{-1} * h_\alpha(\xi, t) \times h_\alpha(\xi, t)^{\frac{b}{1-b}}] \quad (7)$$

where $\delta_\alpha^{-1} * f$ indicates a deconvolution. The nice thing about this approach is that at equilibrium the free energy estimate simply states $A = A + V_b - V_b$ which is free from artifacts in the bias potential.

In practice one may collect a histogram on the same domain as h_α , called h_o , and take the approximation $h_o \approx \delta_\alpha^{-1} * h_\alpha$ into equation (7). This histogram may be collected as a standard histogram or it may be collected the same way as h_α but with a very small Gaussian width. At equilibrium $h_o \propto e^{-\beta\mu}$ so again the free energy estimate is free of artifacts caused by the bias and one does not need to wait for artifacts to be healed.

In Ref. 1 we also noted that equation (3) maps onto the biasing potential used to motivate well-tempered metadynamics given the following parameter conversions,

$$\begin{aligned} \omega &= \beta^{-1}cb \\ \Delta T &= \beta^{-1}b/(1-b) \end{aligned} \quad (8)$$

Plugging these into equation (3) gives

$$V_b(\xi, t) = \Delta T \ln\left(\frac{\omega}{\Delta T}h_\alpha(\xi, t) + 1\right) \quad (9)$$

which is equation 2 in the introduction of well-tempered metadynamics, Ref. 2.

Plugging the WTmetaD parameters into equation (6) gives

$$A(\xi, t) = -(1 + \frac{\beta^{-1}}{\Delta T})V_b(\xi, t) \quad (10)$$

which is exactly the WTmetaD rule for computing free energy from the bias potential.

The WTmetaD bias potential update is motivated by considering the time derivative

$$\begin{aligned} \dot{V}_b(\xi, t) &= \frac{\omega\delta_\alpha(\xi(\mathbf{x}_t) - \xi)}{(\frac{\omega}{\Delta T}h_\alpha(\xi, t) + 1)} \\ &= \omega e^{-V_b(\xi, t)/\Delta T} \delta_\alpha(\xi(\mathbf{x}_t) - \xi) \end{aligned} \quad (11)$$

This is equation 3 of Ref. 2 and defines the “hill height” to be $e^{-V_b(\xi, t)/\Delta T}$.

Based on equation (11), we can define a bias update with the finite difference

$$\dot{V}_b(\xi, t + dt) = \frac{1}{dt}(V_b(\xi, t + dt) - V_b(\xi, t)) \quad (12)$$

to obtain

$$\begin{aligned} V_b(\xi, t + dt) &= V_b(\xi, t) + \\ &dt \omega e^{-V_b(\xi, t+dt)/\Delta T} \delta_\alpha(\xi_{t+dt} - \xi) \end{aligned} \quad (13)$$

where $\xi_{t+dt} \equiv \xi(\mathbf{x}_{t+dt})$ is the center of the most newly deposited hill. This is the update consistent with equation (9), and is simply a numerical solution to equation (11).

It can be said that this update “drops” the Gaussians δ_α along the trajectory, where the Gaussian heights are scaled by $e^{-V_b(\xi, t+dt)/\Delta T}$. Notice that if the effects of finite α are ignored then it is straight forward to conclude that the limiting equilibrium distribution of ξ generated by equation (13) is $e^{-V_b/\Delta T}e^{-\beta(A+V_b)}$. This distribution is uniform on Ω exactly when equation (10) is satisfied. Thus, we concluded in Ref. 1 that results pertaining to mABP could be extended to WTmetaD, generating what we called a “parameter-free” approach. The idea was that equation (7) can correct for α -related errors which makes the method exact for any parameter choice provided sufficient sampling. It was known at the time that large hills caused deteriorated convergence of metadynamics.[6]

However, a distinction of methods can be made by considering the *implemented* update for WTmetaD

$$\begin{aligned} V_b^{WTM}(\xi, t + dt) &= V_b^{WTM}(\xi, t) + \\ &dt \omega e^{-V_b^{WTM}(\xi_{t+dt}, t)/\Delta T} \delta_\alpha(\xi_{t+dt} - \xi) \end{aligned} \quad (14)$$

which is equation 1 of Ref. 4. In Ref. 2 “s” is written for ξ and “s(t)” is written for ξ_{t+dt} . Notice that in Ref. 2 “s(t)” does not appear in the hill height. To obtain the update in equation (14) from the update in equation (13), ξ_{t+dt} must be substituted for ξ in the hill height and nowhere else. The consequences of this substitution have never been formally addressed and are not always clearly implied by notation convention in the literature. This slight notational difference actually distinguishes two very different biasing potentials.

WTmetaD can be obtained from mABP under the exchange

$$e^{-V_b(\xi, t+dt)/\Delta T} \Rightarrow e^{-V_b(\xi_{t+dt}, t)/\Delta T} \quad (15)$$

which redefines the hill height. In reference 1 we showed how to shift the point of evaluation in exponential factors like these and, borrowing those results, this exchange remains faithful to mABP and the motivating expression (11) when α is small so that mollification artifacts do not enter the bias potential. The differences between these bias potentials only emerges when α is large and is summarized by two observations. First of all, the convergence time of WTmetaD begins to suffer since artifacts must be healed prior to convergence.[4, 6]. Second, mABP becomes increasingly less efficient as α is increased because the mollification artifacts prevent the bias from flattening the landscape.

The goal of the following sections is to obtain an artifact-independent free energy estimate for metadynamics that is applicable to the *implemented* metadynamics. This would enable the use of larger hill widths, providing better sampling enhancement, and give confidence that artifacts are irrelevant for any particular application. The first step though is to get a better idea of the role of the hill width in determining the long-time limit of the bias, as so far the literature demonstrates that the method converges without artifact while at the same time demonstrating that it is impractical to converge exactly in the face of artifacts.

B. Consistency for tempered metadynamics

It is straightforward to cast WTmetaD in an alternative form by thinking of the update as a time integral and inserting a Dirac delta function. To do this, we partition the time t into n intervals of $dt = t/n$, where $t_1 = 1dt$, $t_2 = 2dt$, $t_n = ndt = t$, and define the weighted histogram

$$\begin{aligned} w(\xi, t) &= \\ \lim_{n \rightarrow \infty} \sum_{i=1}^n e^{-V_b^{\text{WTM}}(\xi(x_{t_i}), t_i - dt)/\Delta T} \delta(\xi - \xi(x_{t_i})) dt & \quad (16) \\ &= \int_0^t e^{-V_b^{\text{WTM}}(\xi_s, s)/\Delta T} \delta(\xi - \xi(x_s)) ds \end{aligned}$$

The WTmetaD bias can now be expressed as

$$V_b^{\text{WTM}}(\xi^*, t) = \omega \int_{\Omega} w(\xi, t) \delta_{\alpha}(\xi - \xi^*) d\xi \quad (17)$$

The standard WTmetaD is recovered by changing integration order, integrating ξ , and discretizing time as shown in equation (16). The bias is now expressed in terms of a histogram, and this histogram affords some leverage towards understanding the long-time limit of the bias and its dependence on α .

With equation (17), it is possible to make a consistency check by assuming that the bias allows the dynamics to reach equilibrium. Suppose that at time τ a WTmetaD simulation has completely converged and that simulation times beyond τ only produce an equilibrium trajectory.

Splitting the histogram as follows

$$\begin{aligned} w(\xi, t) &= w(\xi, \tau) + \\ \int_{\tau}^t e^{-V_b^{\text{WTM}}(\xi(x_s), s)/\Delta T} \delta(\xi - \xi(x_s)) ds & \quad (18) \\ &= w(\xi, \tau) + \rho(\xi, t - \tau) \end{aligned}$$

allows us to write the bias as

$$\begin{aligned} V_b^{\text{WTM}}(\xi, t) &= V_b^{\text{WTM}}(\xi, \tau) + \\ \omega \int_{\Omega} \rho(\xi^*, t - \tau) \delta_{\alpha}(\xi^* - \xi) d\xi^* & \quad (19) \end{aligned}$$

Since we have assumed equilibrium sampling, we expect that $\rho \propto e^{-V_b^{\text{WTM}}/\Delta T} e^{-\beta\mu} e^{-\max[V_b^{\text{WTM}}]}$. This last factor indicates that the drifting zero-of-energy for the bias is sending the hill heights to zero regardless of robust sampling. Equation (19) clearly implies that the system can only come to equilibrium if $\rho * \delta_{\alpha} = 0$, or $\rho * \delta_{\alpha} = \text{const} \times e^{-\max[V_b^{\text{WTM}}]}$, leaving the bias potential conservative. For the non-trivial case, ρ must be constant across Ω . From the definition of μ we see that ρ is constant when $V_b^{\text{WTM}} = -A/(1 + \beta^{-1}/\Delta T)$. Thus, WTmetaD produces an exact free energy without any mollification artifacts assuming convergence.

This conclusion is identical to what we speculated above (while ignoring the finite hill width) and it is identical to the conclusion of Ref. 4. While we arrive at equation (19) via an alternative route, the equation is just a restatement of equation 7 in Ref. 4 and to aid the comparison we have used ρ to indicate the reweighted equilibrium distribution which was written as $e^{-V_b^{\text{WTM}}/\Delta T} \rho_b$ there. ($\rho = e^{-V_b^{\text{WTM}}/\Delta T} \rho_b$) Here, we have assumed convergence and checked what that converged state must be. Dama and co-workers[4] have provided arguments that WTmetaD does converge by further analysing $\dot{V}_b^{\text{WTM}} = \rho * \delta_{\alpha}$.

At this point, knowing that WTmetaD converges without mollification artifact begs the question of how this is achieved in practice, and how this impacts convergence times.

C. Mollification artifacts and convergence time

Equation (17) shows that the bias potential is related to a Gaussian “blurring” of w . If the bias converges exactly to the underlying free energy, then w must take very specific shapes any time α is large compared to the features of the free energy. In particular, a Gaussian filter will reduce the amplitude of features having spatial frequencies that are large compared to α . Since WTmetaD converges to the free energy regardless of α , some means of compensating for this attenuation of high frequencies must be present.

The long-time limit of WTmetaD gives us

$$\frac{-A(\xi^*)}{\omega(1 + \frac{\beta^{-1}}{\Delta T})} = \int_{\Omega} w(\xi, \infty) \delta_{\alpha}(\xi - \xi^*) d\xi \quad (20)$$

after combining equation (17) and the relationship between the free energy and bias potential. Thus, the histogram that is built during WTmetaD simulations, w , must display the high frequency shapes of A with exaggerated amplitude or those shapes will be washed out by the convolution with δ_α .

We can follow this through with an analytic model. While more complex issues may exist, the following example gets at the mechanics of avoiding mollification. We take a one dimensional collective variable on $[-\infty : \infty]$, where the target free energy is

$$A(\xi) = \cos(\nu \xi) \quad (21)$$

The effects of δ_α can be computed exactly,

$$\begin{aligned} & \int_{-\infty}^{\infty} \delta_\alpha(\xi - \xi^*) A(\xi) d\xi \\ &= \alpha \sqrt{\pi} \exp[-(\alpha \nu)^2 / 4] \cos(\nu \xi^*) \end{aligned} \quad (22)$$

showing that the attenuation of frequency (ν) is related to how large α is relative to that frequency. Up to an additive constant, the histogram that WTmetaD must generate can be written as a function of α

$$w(\xi, \infty) = -\exp\left[\frac{\alpha^2 \nu^2}{4}\right] \cos(\nu \xi) / C(\alpha) \quad (23)$$

where $C(\alpha) = \alpha \sqrt{\pi} \omega(1 + \frac{\beta^{-1}}{\Delta T})$. It is straight forward to check that with this histogram $-V_b^{WTM}(\xi, \infty)(1 + \frac{k_B T}{\Delta T}) = A$.

When $\exp[\alpha^2 \nu^2] \approx 1$ WTmetaD converges with $w(\xi, t) = -\cos(\nu \xi) / C(\alpha)$, up to some additive constant. As α increases, the number of samples required to converge increases as $\exp[\alpha^2 \nu^2]$ for each ξ . The histogram w can only grow logarithmically with time due to the tempering, yet the density of samples required to converge increases exponentially with α . Thus, the WTmetaD convergence time must be of order $\exp[\alpha^2 \nu^2]$ for this model system. This explains why slow convergence is observed for WTmetaD when α is large[6]. (Notice that because of $C(\alpha)$ this analysis also suggests that convergence time increases for small α because $1/C(\alpha)$ becomes large.)

This analysis suggests an imbalance in objectives: The number of samples on Ω that are required to converge are increasing rapidly with α while the importance of the samples is being diminished by that same accumulation. It becomes more and more costly to heal the artifacts as samples are collected, eventually becoming impractical to observe convergence of the free energy estimate.

In the next section we consider alternate tempering schedules with the idea of improving the rate of convergence by either increasing the speed at which w evolves or by adjusting the target free energy estimate so that mollification does not scale the convergence time.

D. Generalized tempering

Consider generalizing the weighted histogram w as

$$w(\xi, t) = \int_0^t f(\xi, s) \delta(\xi - \xi(x_s)) ds \quad (24)$$

where the bias is again given by equation (17). If we take $f(\xi, s) = 1$ we have standard metadynamics. When $f = 1$ the bias must converge to $-A$ if $\rho * \delta_\alpha = \text{const}$. Of course, because f is a constant, the metadynamics bias never reaches this state. By always adding fixed height hills, metadynamics ensures that $\rho * \delta_\alpha$ cannot be a constant across Ω . This point is appreciated as the source of the oscillations in the metadynamics free energy estimate and demonstrates that a consistency check is not a proof.

If we take $f(\xi, s) = e^{-V_b^{WMT}(\xi, s)/\Delta T}$ we have WTmetaD. Additionally, we can take $f(\xi, s) = f(s)$ where $f(s)$ is a strictly decaying function of time. One special choice of $f(s)$ has been given in TTmetaD (transition-tempered metadynamics[7]). The bias potential generated with the TTmetaD update converges to $-A$ but, unlike metadynamics, leads to asymptotic convergence because the updates are vanishing. TTmetaD does suffer from the problem of artifacts, but will converge eventually.

Any given tempering can be implemented as

$$\begin{aligned} V_b^{TM}(\xi, t + dt) &= V_b^{TM}(\xi, t) + \\ &dt \omega f(\xi_{t+dt}, t) \delta_\alpha(\xi_{t+dt} - \xi) \end{aligned} \quad (25)$$

and the free energy can be estimated as

$$A(\xi, t) = \beta^{-1} \ln \left(f(\xi, t) \right) - V_b^{TM}(\xi, t) \quad (26)$$

which follows from our consistency check. We write TM just to mean any ‘‘Tempered Metadynamics.’’ One can check that this reduces to the WTmetaD rule upon substitution of the well-tempered schedule $f = e^{-V_b^{TM}/\Delta T}$.

Notice that if $f(\xi, t) \propto e^{\beta \mu}$ then equation (26) would capture the spirit of the mABP free energy estimate. Plugging this f into equation (26) gives $A = \mu + V_b^{TM} - V_b^{TM}$. To achieve this, one may collect the histogram h_\circ as defined earlier (also subject to $h_\circ(\xi, 0) = 0$) and use the tempering schedule

$$f(\xi, t) = \frac{1}{rh_\circ(\xi, t) + 1} \times \left(\max_{\xi \in \Omega} [rh_\circ(\xi, t) + 1] \right)^m \quad (27)$$

which has been designed after equation (11). Based on the fact that at long times $h_\circ \propto e^{-\beta \mu}$, at equilibrium equation (26) writes $A = \mu - V_b^{TM}$ which is exactly what we want. This tempering serves to simultaneously shrink the hills and reweight the free energy estimate. Setting $0 \leq m < 0.5$ allows some control over the decay rate of the hills while still making sure the hills eventually vanish. Recently Fort and co-workers looked at importance

sampling and suggested that this sort of decay rate (*i.e.* t^{m-1}) may be an improvement for importance sampling and tempered metadynamics[8]. However, after numerical experiments below we suggest that $m = 0$ is always the best option as taking larger values only delays equilibration of this scheme.

The parameter r plays the role of $\omega/\Delta T$ in WTmetaD or $c(1-b)$ in mABP. The meaning of ΔT is lost however, so we simply use $0 \leq r$. Indeed, the numerical results below indicate that $r = \omega/\Delta T$ is not an optimal choice in general. The smaller r , the slower the decay of the hills and when $r = 0$ standard metadynamics is recovered. Since h_o is defined as

$$h_o(\xi, t) = \int_0^t \delta_o(\xi(x_s) - \xi) ds$$

r has units of dt^{-1} . Ultimately r provides a more fine control over the hill decay than does m because it avoids changing the envelope of the hill decay.

The only concerns for this scheme are that (1) the hills do not decay too fast, leaving the bias to only inefficiently push trajectories over high barriers, and (2) that the hills decay fast enough to let the dynamics equilibrate on a reasonable timescale. The same two concerns exist when specifying any tempered metadynamics parameter set.

We refer to the tempering schedule of equation (27) as “ μ -tempered” metadynamics or μ TmetaD because it uses the PMF of the biased ensemble (μ) to temper. This scheme bypasses the mechanism for healing artifacts and leans on the shrinking hill heights to guarantee that equilibrium will be met. Once in equilibrium the free energy estimate reduces to an identity $A = A$ (assuming Ω is robustly sampled), which was designed in analogy to the mABP free energy estimate.

Next, we use numerical experiments to compare WTmetaD and μ TmetaD.

E. Numerical experiments

First we demonstrate the role of w in WTmetaD. We have implemented WTmetaD according to equation (14) so that we may compare WTmetaD simulations to our expectations for the histogram w . We used $b = 0.8, c = 0.01, k_B T = 0.2$ (with no experimentation) and the Langevin integrator from reference 9. The target free energy is given by equation (21). We took ξ on the interval $[-1 : 1]$ and set $\nu = 5\pi$ for this numerical example. We discretized Ω into 200 bins.

The histogram w is not normally tracked during WTmetaD simulations but can easily be computed by collecting a standard histogram on Ω that is reweighted by the hill height scaling. That is, each time a bin on Ω is visited, the histogram is updated by adding $e^{-V_b^{WTM}/\Delta T}$ to that bin.

In figure 1, we plot equation (23) and the histogram collected from a standard WTmetaD simulation implemented as equation (14), called \tilde{w} . As expected, when

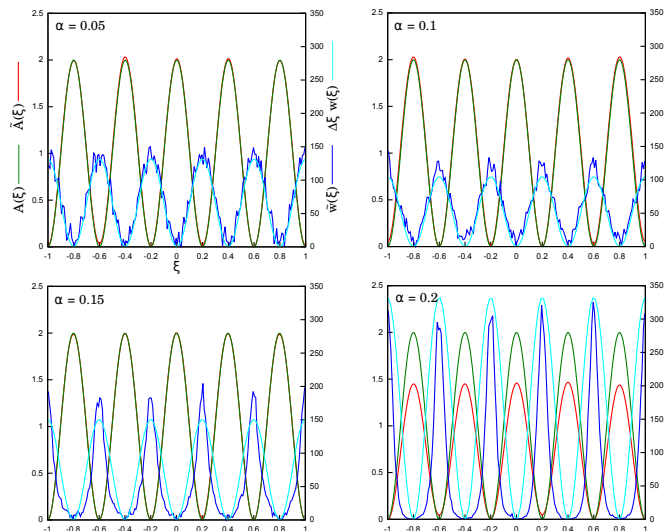


FIG. 1: Here we demonstrate the impacts of large α and show how w must exaggerate the shape of A to account for the convolution in equation (17). \tilde{A} and \tilde{w} are the WTmetaD estimates. $\Delta\xi$ is the bin width used in simulations and appears for dimensional consistency. One million MD steps were used.

α increases the high frequency amplitudes of A must be exaggerated in the WTmetaD estimate of w . Comparing $\alpha = 0.05$ with $\alpha = 0.2$, we see that the peaks in \tilde{w} are twice as high at the larger α yet the barriers of the free energy are still under estimated.

To look at how μ TmetaD behaves in practice, we considered the case of alanine dipeptide that was used to demonstrate poor convergence for WTmetaD at large values of α in Ref. 6. In particular it was demonstrated that α can introduce artifacts that are difficult to heal, thus motivating a free energy estimate alternative to equation (10) which would also allow the use of more aggressive hill functions. The adaptive Gaussian schemes themselves are unable to improve upon the WTmetaD free energy estimate without abandonment of equation (10).

Tempered metadynamics, mABP, and a “SHUS” scheme based on Ref. [8] (see appendix) were implemented in GROMACS-4.5.5[10] and the $\phi - \psi$ dihedral space was broken into a 300-by-300 bin grid. Using CHARMM27 and TIP3P water, Alanine dipeptide was solvated and equilibrated for three nanoseconds in the NPT ensemble. Free energy calculations were done in the NVT ensemble. All codes, inputs, and instructions to reproduce these simulations can be obtained from the ABPenabledGROMACS[11] github page.

The biasing parameters were set to $b = 0.8$ and $c dt = 0.1$. ($\omega dt = \beta^{-1} c dt b$, and $\Delta T = \beta^{-1} b / (1 - b)$ where dt is the timestep.) Thus, $\omega dt = 0.1992$ and $\Delta T = 9.96$ where both are reported in kJ/mol. A reference free energy A_{ref} , was computed from a 100 nanosecond simulation that was biased with a static bias potential. The reference free energy was obtained via equilibrium reweighting

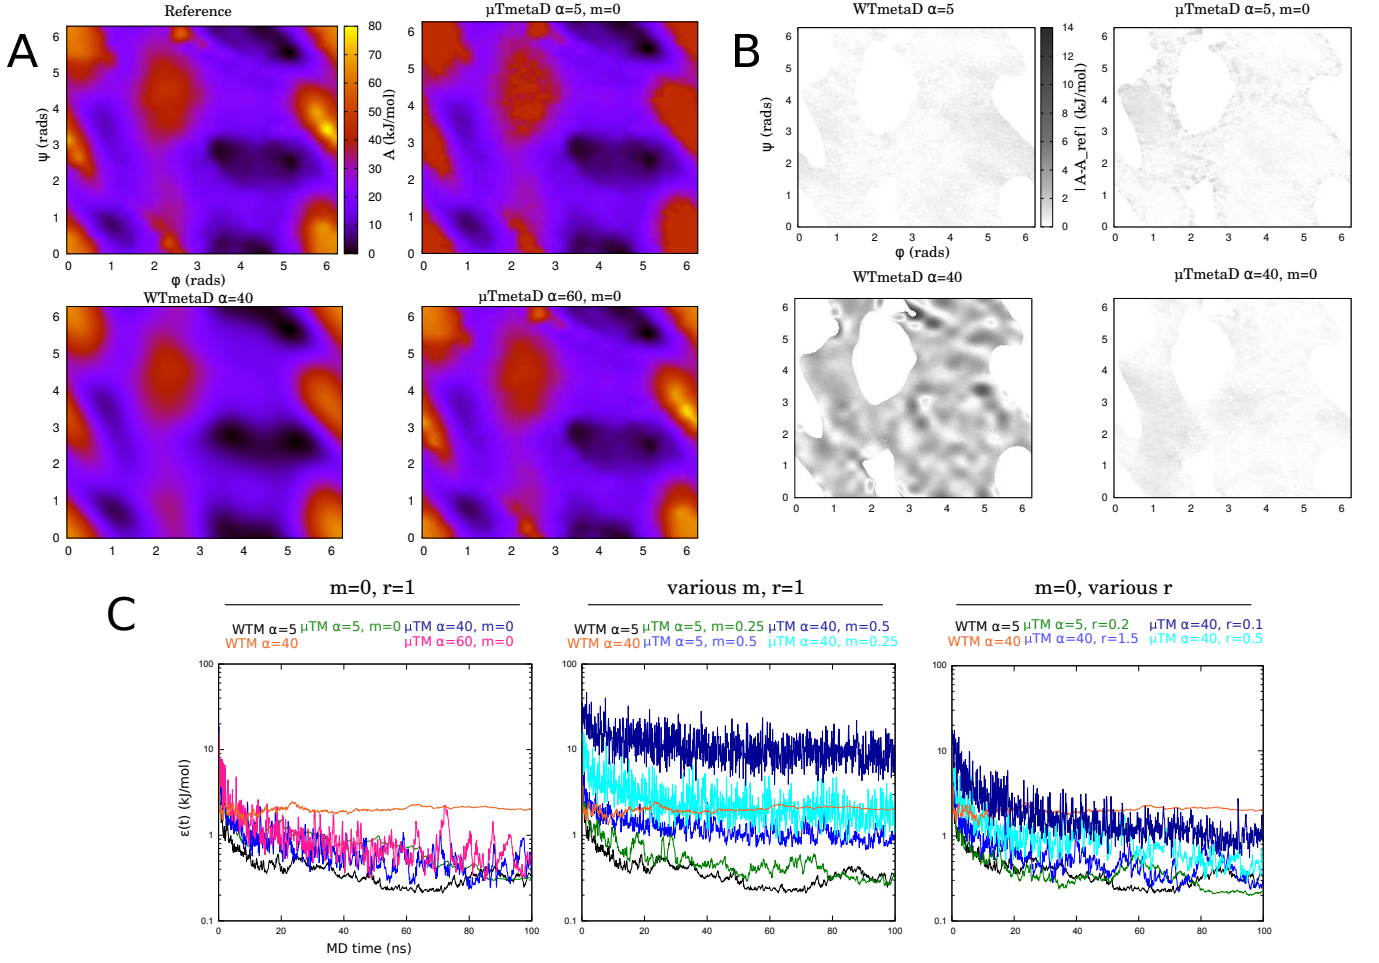


FIG. 2: (A) The reference free energy is shown along with estimates from the different temperings. (B) The point by point difference between the free energy reference and an estimate for regions contributing to $\epsilon(t)$. (C) The value of $\epsilon(t)$ for the different temperings, and for different r and m values. Rather than r we actually report the unitless value $r \times dt$.

ing. The static bias potential itself was computed from a 10 nanosecond mABP simulation at $\alpha = 8$ degrees. The level of convergence was then defined as

$$\epsilon(t) = \frac{1}{N} \sum_{i,j} |\hat{A}(\xi_{i,j}, t) - A_{\text{ref}}(\xi_{i,j})| \times \Theta(30 - A_{\text{ref}}(\xi_{i,j})) \quad (28)$$

where the free energy estimate \hat{A} may be derived from μ TmetaD using equation (26) with $h_o = h_{2^\circ}$ or WTmetaD using equation (10). Θ is the Heaviside step function and the number N is the number of grid points where $A_{\text{ref}} < 30$ kJ/mol. Reference 6 defined convergence in a similar way using nearly the same energy cut-off. Simulations were run with $\alpha = 5, 40$ degrees for 100 nanoseconds each. The choice $\alpha = 40$ was taken from Ref. 6. All the simulation results are summarized in figure 2.

In panel A of figure 2 we show the reference free energy and estimates from different temperings. The estimate from μ TmetaD with $r = 1$ and $\alpha = 5$ clearly shows that the sampling is inefficient in the high energy regions of

the landscape. The “energy volume” of each bias update is too small. In this limit the bias updates do more to diminish the hill heights than they do to enhance sampling. The analogous case is setting ΔT too low in a WTmetaD simulation. In panel C we show that with $r = 0.2$ the convergence of μ -TmetaD matches that of WTmetaD when $\alpha = 5$. The smaller r allows the hills to decay more slowly but maintains the $1/t$ decay.

Free energy estimates from WTmetaD with $\alpha = 40$ and μ TmetaD with $\alpha = 60$ and $r = 1$ are also shown in panel A of figure 2. The smoothing effects of α are visible by eye in the WTmetaD estimate, yet the μ TmetaD estimate is as sharp as the reference even though this simulation used a value of α that is significantly larger. Also notice the improved sampling of high energy regions for the larger α , reflecting the larger filling-rate that motivates using wider hills.

In panel B of figure 2 we show the point by point difference in free energy that is used to evaluate $\epsilon(t)$. When α is small, WTmetaD produces a sharp free energy estimate with little error. When $\alpha = 40$, WTmetaD pro-

duces a far less crisp estimate resulting in point by point errors that can be significant (i.e., up to 5 or 6 $k_B T$). The point by point error in the μ TmetaD result at $\alpha = 40$ is, in contrast, still as small as it was when $\alpha = 5$. This demonstrates that the free energy estimate is free from α -related artifacts, where all simulations have evolved for the same duration (100ns).

Finally, in panel C of figure 2 we show plots of $\epsilon(t)$ for the different temperings and different parameters. Keep in mind that the values of ϵ are smaller than those shown in panel B simply because ϵ is an average. What is clear from panel C is that μ TmetaD is converging rapidly for $\alpha = 5, 40, 60$ (with $m = 0, r = 1$) while WTmetaD is reaching a plateau for larger α . This plateau is consistent with the visible differences shown in panel A and the point by point error in panel B, as well as with the previous reports of unobserved convergence at $\alpha = 40$. Clearly, μ TmetaD displays artifact independent convergence.

Panel C of figure 2 also shows the behavior of μ TmetaD for $m = 0.25, 0.5$ when $r = 1$. We see that increasing m can impede convergence no matter the size of α . The μ TmetaD free energy estimate is only meaningful at equilibrium, where equation (26) can be interpreted as $A + V_b^{TM} - V_b^{TM}$. Taking $m > 0$ delays equilibration and clearly has negative impact on convergence time. We observed the same qualitative behavior with the scheme in Appendix equation (C1) (results not shown, but the scheme is available in the github[11] repository). Our summary is that taking decay rates that go as t^{m-1} does not improve convergence and instead only slows equilibration leading to higher levels of noise in the free energy estimate. The source of error in the μ TmetaD free energy estimates stems from uncertainty in the depths of the basins and not in the resolution of the estimate. Since the hills decay very slowly, the noisy behavior of standard metadynamics is partly recovered.

Lastly, panel C shows that r can be used to change the hill decay without the drastic delay of equilibration that is caused by changing m . The errors that remain in the free energy when $\alpha = 40, m = 0$, and $r = 0.1$ stem from incorrect determination of relative depths of the free energy minima. The free energy estimate itself is as crisp as what is shown in panel A when $\alpha = 60$. The tempering is much more robust to adjustments in r than it is to m , suggesting $m = 0$ is a best choice. Since r simply needs to be a function of time or a constant, a wide range of variations can be imagined and we do not try to find optimal solutions here.

III. CONCLUSION

The difference between mABP and WTmetaD traces back to a single approximation (equation (15)) which causes the methods to diverge from one another when α is large. The two main consequences are that the results of mABP do not apply to metadynamics and that metadynamics must heal all artifacts in the bias potential

before the free energy can be accurately estimated.

Because the hills of WTmetaD are tempered and shrinking, the process of healing artifacts extends convergence times considerably. The artifacts themselves require more samples accumulated to w and all tempering schedules decrease the weight of each MD timestep when it contributes to w . In analogy with mABP, a tempering scheme that does not need to heal artifacts to reach an accurate free energy estimate was proposed, called μ -tempered metadynamics. This scheme was compared to WTmetaD in the common benchmark system of solvated alanine dipeptide.

The μ TmetaD combines the spirit of mABP with the tempering of metadynamics to achieve a method that may represent the best of both worlds: Convergence times that do not scale with α , and full exploitation of wide hills. In mABP, configuration space exploration begins to suffer in the limit of large α . [1] In traditional metadynamics schemes configuration space exploration is improved by large α but the free energy estimate (and convergence) suffers. μ TmetaD allows wide hills to be used for the best sampling efficiencies and produces artifact independent convergence times.

IV. ACKNOWLEDGEMENTS

All computations were performed on the Van Andel Research Institute compute cluster.

Appendix A: Pseudo codes

Here we will write simple pseudo-code for each adaptive biasing scheme. Each scheme stores information on a grid, which is used to construct the biasing potential and biasing force. To stress the similarity of these methods we define \mathcal{M} and \mathcal{M}' as two arrays stored in memory. We store information for the bias potential in \mathcal{M} and information for the biasing force in \mathcal{M}' .

WTmetaD and μ -TmetaD

Suppose it is time to deposit a new hill. The new hill has center $\xi^* = \xi(x_t)$, and the hill height scale, s , of this hill is

$$f(\xi^*, t) = \omega \exp[-\mathcal{M}(\xi^*, t)/\Delta T]$$

The arrays \mathcal{M} and \mathcal{M}' are updated at all grid points ξ as

$$\begin{aligned} \mathcal{M}(\xi, t + dt) &= \mathcal{M}(\xi, t) + f \delta_\alpha(\xi^* - \xi) \\ \mathcal{M}'(\xi, t + dt) &= \mathcal{M}'(\xi, t) + f \delta'_\alpha(\xi^* - \xi) \end{aligned} \quad (A1)$$

where

$$\delta'_\alpha(\xi^* - \xi) = \frac{2}{\alpha^2}(\xi^* - \xi)\delta_\alpha(\xi^* - \xi)$$

The biasing force at ξ is given by $-\mathcal{M}'(\xi, t)$. The bias potential is \mathcal{M} . For periodic collective variables, the minimum image convention must be observed. We take the

traditional factor of dt into ω (and the mABP parameter c) here.

μ TmetaD is the same as WTmetaD except that

$$f(\xi^*, t) = \frac{\omega}{rh_o(\xi^*, t) + 1} \left(\max_{\xi \in \Omega} [rh_o(\xi, t) + 1] \right)^m$$

where h_o is a histogram as outlined above.

mABP

Again let $\xi^* = \xi(x_t)$ denote the position of the trajectory in this update. The arrays stored in memory are updated at all grid points ξ as

$$\begin{aligned} \mathcal{M}(\xi, \tau + dt) &= \mathcal{M}(\xi, \tau) + \delta_\alpha(\xi^* - \xi) \\ \mathcal{M}'(\xi, \tau + dt) &= \mathcal{M}'(\xi, \tau) + \delta'_\alpha(\xi^* - \xi) \end{aligned} \quad (\text{A2})$$

The biasing force at ξ is given by

$$\frac{-cb\beta^{-1}\mathcal{M}'(\xi, \tau)}{c(1-b)\mathcal{M}(\xi, \tau) + 1}$$

and the bias potential is $\beta^{-1} \frac{b}{1-b} \ln[c(1-b)\mathcal{M} + 1]$ but only the biasing force is evaluated during simulations.

Ideally δ_α is pre-computed at the beginning of simulation and stored for look-up. All simulations in this work used pre-computed δ_α .

Appendix B: mABP as a tempering schedule

For completeness we note that frequency attenuation may be introduced into the bias potential with the following definition for f ,

$$w(\xi, \xi^*, t) = \int_0^t e^{-V_b(\xi^*, s)/\Delta T} \delta(\xi - \xi(x_s)) ds \quad (\text{B1})$$

Plugging into equation (17), the bias becomes

$$\begin{aligned} V_b(\xi^*, t) &= \omega \int_0^t \frac{\delta_\alpha(\xi(x_s) - \xi^*)}{\frac{\omega}{\Delta T} \int_0^s \delta_\alpha(\xi(x_\tau) - \xi^*) d\tau + 1} ds \\ &= \Delta T \ln \left(\frac{\omega}{\Delta T} h_\alpha(\xi^*, t) + 1 \right) \end{aligned} \quad (\text{B2})$$

which recovers mABP. In passing from the first line to the second line, we solved the differential equation given in equation (11). The bias potential is now a log-histogram so updates to the bias potential decay as $1/t$ even though all observations of ξ contribute equally to h_α . This $1/t$ decay rate motivated the exponential temperings used in WTmetaD and TTmetaD[7], but requires no tempering or reweighting when the bias is defined directly as a log-histogram.

Because the solution to the differential equation in equation (11) is unique given a boundary condition $V_b(\Omega, 0) = 0$, one may conclude that the only scheme consistent with equation (11) is mABP. The WTmetaD update in equation (14) is not strictly a numerical solution to equation (11). This observation does not change the fact that WTmetaD is convergent, and rather argues that equation (11) is not useful for describing WTmetaD (or other metadynamics).

Appendix C: A SHUS-like scheme

Taking advantage of the flexibility in equation (24), we proposed a tempering inspired by “Algorithm 3” in reference 8. In particular,

$$f(\xi, t) = \frac{1}{\ln[V_b^{TM}(\xi, t) + 1]^{\frac{\gamma}{1-\gamma}} + 1} \quad (\text{C1})$$

although this is not identical to “Algorithm 3” in Ref. 8. Note that the parameter value inputs for the code on the github page[11] is the value of $\gamma/(1-\gamma)$.

-
- [1] B. M. Dickson, Phys. Rev. E **84**, 037701 (2011).
 - [2] A. Barducci, G. Bussi, and M. Parrinello, Phys. Rev. Lett. **100**, 020603 (2008).
 - [3] A. Laio and M. Parrinello, PNAS **99**, 12562 (2002).
 - [4] J. F. Dama, M. Parrinello, and G. A. Voth, Physical Review Letters **112**, 240602 (2014).
 - [5] A. Laio, A. Rodriguez-Fortea, F. L. Gervasio, M. Ceccarelli, and M. Parrinello, J. Phys. Chem. B **109**, 6714 (2005).
 - [6] D. Branduardi, G. Bussi, and M. Parrinello, Journal of Chemical Theory and Computation **8**, 2247 (2012).
 - [7] J. F. Dama, G. Rotskoff, M. Parrinello, and G. A. Voth, Journal of Chemical Theory and Computation (2014).
 - [8] G. Fort, B. Jourdain, T. Lelievre, and G. Stoltz, arXiv

- preprint arXiv:1410.2109 (2014).
- [9] M. Allen and D. Tildesley, *Computer simulation of liquids* (New York: Oxford, 1989).
- [10] B. Hess, C. Kutzner, D. van der Spoel, and E. Lindahl, J. Chem. Theory Comput. **4**, 435 (2008).
- [11] B. Dickson, <https://github.com/BradleyDickson/ABPenabledGROMACS> (2015), accessed: 2015-06-13.